Dynamic Mapping for Multiview Autostereoscopic Displays

Jing Liu[a], Tom Malzbender[b], Siyang Qin[a], Bipeng Zhang[a], Che-An Wu[a] and James Davis[a]

[a]UC Santa Cruz, 1156 High St, Santa Cruz, United States;
[b]Cultural Heritage Imaging, 2325 3rd Street, San Francisco, United States

ABSTRACT

Multiview autostereoscopic displays have several image artifacts which prevent widespread adoption. Crosstalk between adjacent views is often severe, stereo inversion occurs at some head positions, and legacy 2-view content is difficult to display correctly. We introduce a method for driving multiview displays, dynamically assigning views to hardware display zones, based on potentially multiple observer's current head positions. Rather than using a static one-to-one mapping of views to zones, the mapping is updated in real time, with some views replicated on multiple zones, and some zones left blank. Quantitative and visual evaluation demonstrates that this method substantially reduces crosstalk.

Keywords: Autostereoscopic, Crosstalk, Eye Tracking, Dynamic Mapping

1. INTRODUCTION

Stereoscopic display technology provides a 3D viewing experience, giving a closer reproduction of the physical world. 3D displays have had important impacts in scientific visualization, engineering, and entertainment. Glasses-free 'Autostereoscopic' 3D displays produce 3D scenes without requiring viewers to wear stereo glasses. The display presents different views to observers in a discrete set of spatial zones. However, user experience in commercially available autostereoscopic displays suffers from three main issues. First, the display may exhibit significant crosstalk, where an observer in a given zone will also see contributions from the views intended for neighboring zones. Crosstalk is a widely recognized problem in multiview autostereoscopic research and can cause visual discomfort for viewers watching such a display.[1] Second, the observer experiences stereo inversion in positions which place the left and right eyes in neighboring repetitions of the display zones. This forces the user to adjust their head position until the proper 3D effect is apparent. Third, multiview displays are not directly compatible with the vast majority of existing stereo 3D content, which is recorded and stored in a 2-view format. The most common solution is to simply map the two available views onto alternating display zones, but this exacerbates the crosstalk and inversion issues above.

In this paper we propose software dynamic mapping of views to physical display zones. Rather than using a predetermined and static one-to-one mapping of views to zones, views are dynamically allocated to zones based on current eye positions. By replicating some views in multiple zones, and leaving some zones blank, crosstalk can be greatly reduced. In addition, dynamic allocation of views to zones allows stereo inversion artifacts to be eliminated and 2-view content to be displayed comfortably to viewers without glasses.

We have implemented a prototype system using a commercially available 8-zone lenticular display. We measure and compare crosstalk between the original display and when using dynamic mapping. We find crosstalk to be substantially reduced using our method. The cost of these improvements is a requirement for head position tracking, which we implement with a 3D camera and standard computer vision, for multiple viewers.

The primary contribution of this paper is a software-based approach to improve image quality on multiview autostereoscopic displays. We support this contribution with a prototype, and both quantitative and visual evaluation.

Further author information: (Send correspondence to Jing Liu)
Jing Liu : E-mail: jingliu@soe.ucsc.edu, Telephone: 1 831 419 7687

2. RELATED WORK

Multiview display hardware has been well studied and designs exist using lenticular screens,[2] parallax barriers,[3] and multiple projectors.[4] Implementations might have as many as 60 display zones.[5] Good surveys exist,[6–8] and advanced hardware technologies continue to be invented.[9–11] Our work focuses on using software to ameliorate some of the deficiencies common to many of the most common hardware platforms.

Eye tracking has been used to improve the functioning of hardware, often mechanically moving optical elements to steer display zones to the current eye position. In ATTEST, lenticular screens with two views are mechanically adjusted according to the viewer's position.[12] Project MUTED and HELIUM3D replace a conventional backlight with novel steering optics and a horizontally scanned light valve so that zones can be directed to appropriate viewers' eyes.[13] Liou et al. use a synchro-signal LED scanning backlight to reduce crosstalk.[14] Stolle et al. also introduce an electronically steerable backlight.[15] These systems all use head tracking to directly manipulate hardware components, dynamically redirecting zones to a moving viewer.

Our use of eye tracking differs fundamentally from above techniques. Our method uses eye tracking to dynamically map views to the static zones supported by given display hardware. We avoid modifying the physical device, and are largely agnostic to the specific technology employed.

Most similar to our method, Kim et al. describes the dynamic fusion and elmination concept applied to a single eye-tracked viewer.[16,17] Boev et al. proposes a method of optimizing quality and brightness for a single user.[18] Earlier, Woodgate et al. had developed the PIXCON LCD panel, focusing primarily on novel hardware but proposing electronic tracking.[19] However, even though these works all suggest that content swapping based on a viewer's eye position is possible, they lack of strategies for dealing with conflicts occurred when there are multiple viewers. Therefore above systems only demonstrate such application in single viewer scenario. Our work introduces an optimization framework suitable for multiple users, a functioning prototype, and quantitative measurements of results.

The Random Hole Display proposed by Nashel, et al. is in spirit most similar to our work when dealing with conflicts among multiple viewers.[20,21] A customized parallax barrier is utilized to replace aliasing artifacts with high frequency noise. Their method optimizes image quality for multiple viewers by distributing error from conflicting views of individual pixels. We also use an optimization with the goal of distributing error, applied at the zone level for a lenticular display.

Heide et al. develop optimization methods for driving compressive displays that sample lightfields.[22] Their work focuses on a custom hardware design with considerably more flexibility than the commercially available lenticular screens that are the focus of our work. However, our work is related to theirs in that it views the output of the display as a function of a controllable set of inputs, and seeks an optimum solution for that input.

A number of authors have addressed image quality by modifying the content prior to display. Zwicker et al. provide a frequency analysis and prefilter for anti-aliasing.[23] Song-Pei et al. apply shears and stitching to reconstruct displayed light field in a way that removes sharp view transitions.[24] Masia et al. propose a light field retargeting method using optimization.[25] Didyk et al. combine phase based video magnification and antialiasing into a single filtering process.[26,27] With methods that produce a set of greatly improved static views, our work could be used in conjunction to further improve the display quality.

3. OUR METHOD

Autostereoscopic screens: Multiview autostereoscopic displays can be built based on a wide variety of optical principles, however all share a common viewing geometry. The hardware supports a finite number of display zones, each of which directs the screen image in a different angular direction. The screen is driven with a set of imagery rendered or photographed from slightly different views. Under normal operation views are mapped to zones in a static one-to-one manner. The observer's two eyes are in different zones, and thus receive different views, resulting in stereo perception.

Unfortunately the hardware usually supports a limited number of zones, and does not have sharp transitions between them. For example, we use a display that has 8 zones, and relatively broad transitions. This results in substantial crosstalk between views, as shown in Figure 1(a). The average intensity of each view, as seen from a range of horizontal positions is plotted. For example, when a user's eye is placed at position 150mm, this is intended to be zone 5. However in addition to view 5, they will see each of view 4 and view 6 at 40% intensity, resulting in visible double images.

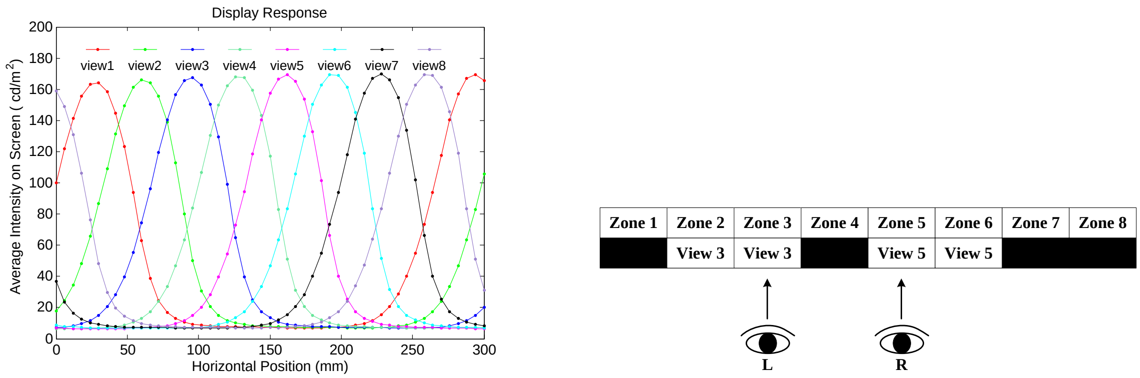

(a) Crosstalk Measurement

(b) Dynamic Mapping Example

Figure 1. (a) Multiview displays often have severe crosstalk. The average visible intensity from each zone was measured at many positions horizontally. Note that all positions have significant contribution from at least two display zones. (b) Our method of dynamic mapping decouples desired views from physical zones. Rendering the same view in several zones and leaving blank zones between eyes can significantly reduce crosstalk. The mapping from views to zones is updated in realtime based on eye position.

An additional complication is that not only do adjacent view zones of multiview displays overlap producing crosstalk, driving the zones themselves is neither independent nor linear. For example, the curves in Figure 1(a) are produced when driving view zones one at a time while measuring screen brightness. However, it is not true that driving Zone1 and Zone2 simultaneously would result in a final intensity which is a linear combination of the values measured for each respective zone. In general, it is necessary to know the optical transfer function of the specific display, whether through calibration or theory.

Dynamic mapping: Crosstalk can be substantially reduced by using eye position as an input and dynamically mapping views to physical display zones. Central to our approach is avoiding crosstalk by ensuring that adjacent view zones that contribute to a single user's eye are driven by weighted copies of an identical image. This simple approach allows the bleedthrough from adjacent physical display zones to still occur, but it does not lead to image degradation. Our method is illustrated in Figure 1(b). In this simple example with a single viewer, we drive both Zone2 and Zone3 with weighted copies of View3. Similarly, it is not necessary to display a view in every zone, and the zone between two eyes can be left black in order to suppress crosstalk.

Multiple viewers may be positioned such that a more complex optimization is needed for determining what image is best displayed in each zone. However the core of our method is the simple observation that our goal is to optimize image quality at the eye positions, not to optimize quality everywhere, and driving the display with something other than a one-to-one mapping of views to zones often results in the minimum error.

Minimizing Error: The error we wish to minimize, and thus our strategy for dynamic mapping must account for two primary goals. First we should minimize crosstalk error. Second, we should account for content sources which have a different number of views than the display.

Whether the source images, $\mathbf{V}(\mathbf{x})$, are virtual or captured from cameras, a change in viewer position, $\mathbf{x}$, induces a parallax change in the image viewed. Since the display device has discrete zones, the *views* are often discretized and denoted as $\mathbf{V} = [V_1, V_2, V_3, \cdots, V_K]$. The hardware device displays different images angularly, in each of a set of N *zones*. The images which are input to the device are denoted $\mathbf{Z} = [Z_1, Z_2, Z_3, \cdots, Z_N]$. Current devices assume that the number of available views, K, is equal to the number of display zones, N. On existing displays views are mapped to zones in a one-to-one fashion, such that $\mathbf{Z} = \mathbf{V}$.

The display device frequently does not closely match the idealized display of discrete zones with sharp boundaries. Neighboring zones exhibit crosstalk, and each display zone influences a wide region of space with maximum intensity in the center of its intended range, and less intensity as the eye position moves into neighboring zones.

The *eye* image, $\mathbf{E}$, actually observed is a function of both the multiple zone images displayed on the device, $\mathbf{Z}$, and the eye position, $\mathbf{x}$. That is, at each position of user space, $\mathbf{E}(\mathbf{x}) = \boldsymbol{\mathcal{E}}(\mathbf{Z}, \mathbf{x})$.

Traditionally there is a direct spatial relationship between views, $\mathbf{V}(\mathbf{x})$, and the desired images at the users eye position, $\mathbf{D}(\mathbf{x})$. However a more complex *view selection policy* is possible. We might, for example, want to ignore spatial relationship to the scene and simply specify that the desired image at the left eye position, $\mathbf{D}(\mathbf{x}_L)$, should always be the left eye view, V_L, in a legacy 2-view movie. Similarly, we might want to detect possible multi-user conflicts in 8-view content and shift the desired views for individual users, even though that would change their virtual viewing direction slightly. In general the view selection policy depends on the application and user preferences for degradation when the device can not provide ideal images. We discuss one possible policy for 2-view content in Section 5.2.

Since the display device has crosstalk and other deficiencies we do not simply drive the closest hardware zones with the desired images, $[\mathbf{D}(\mathbf{x}_1), \mathbf{D}(\mathbf{x}_2), \cdots]$. In general, some other set of images will produce actual observed images which are closer to those that are desired. We refer to this selection of what to drive the device with as the *display algorithm.* Perceptually based metrics for evaluating the difference between two images exist.[28] In this work, we obtain acceptable results by simply minimizing the squared image intensity error between what is desired, and what the device actually produces.

$$\underset{\mathbf{Z}}{\operatorname{argmin}} \sum_{m=1}^{M} [\boldsymbol{\mathcal{E}}(\mathbf{Z}, \mathbf{x}_m) - \mathbf{D}(\mathbf{x}_m)]^2 \tag{1}$$

Display Algorithm: Global optimization over all possible zone images would be prohibitively slow. Our display requires eight input images each with millions of independent pixels. Fortunately we don't need to find an optimum solution, only one that produces low errors, so we use heuristics to prune the search space, and optimize in the more constrained space.

First we assume that the only reasonable value for each input zone image $Z_1, \cdots, Z_N$ is a weighted linear combination of the the desired images at all of the eye locations $\mathbf{D}(\mathbf{x}_1), \cdots, \mathbf{D}(\mathbf{x}_M)$. This corresponds to assuming we should drive the display with things we want to see, not some other image entirely.

$$Z_i = w_{i1} \cdot \mathbf{D}(\mathbf{x}_1) + w_{i2} \cdot \mathbf{D}(\mathbf{x}_2) + \cdots + w_{iM} \cdot \mathbf{D}(\mathbf{x}_M) \tag{2}$$

Since there are N display zones and M eye positions, this assumption defines a weight matrix, $\mathbf{w}$, with $N * M$ terms which encodes the possible solution space.

To further reduce the dimensionality of the search space, for each zone, Z_i, we check to see whether it contributes less than 15% of the full energy to each eye position $\mathbf{x}_j$, that is $max(\mathbf{E_i}(\mathbf{x_j})) < 15\%$. If the contribution is small then we set $w_{ij} = 0$ in $\mathbf{w}$.

We now substitute our restricted definition of $\mathbf{Z}$ from Equation 2 into the general minimization defined in Equation 1, and minimize over the elements of matrix $\mathbf{w}$ which were not set to 0.

$$\underset{\mathbf{w}}{\operatorname{argmin}} \sum_{m=1}^{M} [\boldsymbol{\mathcal{E}}(\mathbf{Z}, \mathbf{x}_m) - \mathbf{D}(\mathbf{x}_m)]^2 \tag{3}$$

4. IMPLEMENTATION DETAILS

Experimental Setup: Our experiments were conducted on a lenticular-based multiview autostereoscopic 3DTV made by Alioscopy. The display uses a slanted lens array as an optical filter, affixed to a normal LCD screen to distribute separate images in each direction. The hardware has 8 horizontal display zones defined. The manufacturer has calibrated the device for optimum performance at 140cm away from the display, and we place our viewing couch at approximately this distance. With this positioning, the display has zones spaced approximately 32mm apart. The average interocular distance in adults is 63mm, corresponding to a two zone separation between eyes on this display.[29] However the statistical range of ocular separation is 50-75mm, and some viewers will thus occasionally have eyes in neighboring zones. Eye positions are tracked by either a pair of webcams or a Kinect 3D camera on top of the display and the information is used for dynamic mapping of views in realtime. Eye tracking is well studied, with surveys available.[30] Our eye tracker with webcams is implemented with a decision cascade of Haar basis functions.[31] We observe the Kinect based tracker to be more robust with real users but either provides acceptable performance.

Calibration: Autostereoscopic display devices typically have optical arrangements more complicated than standard 2D displays. The crosstalk between zones, as well as any spatial or radiometric nonlinearities must be taken into account. These issues are encapsulated in the display transfer function, $\boldsymbol{\mathcal{E}}(\mathbf{Z}, \mathbf{x})$, which describes what the eye actually sees at position, $\mathbf{x}$, when the display zones are driven with a set of images, $\mathbf{Z}$.

We sample the space of both $\mathbf{Z}$ and $\mathbf{x}$, recording the actual observed intensities at all screen pixels using a camera. We store the samples and evaluate $\boldsymbol{\mathcal{E}}(\mathbf{Z}, \mathbf{x})$ by linear interpolation. We have found this simple method empirically sufficient. We have sampled $\mathbf{x}$ as densely as 6mm and as coursely as 32mm, and $\mathbf{Z}$ with as few as 9 samples and as many as 729 samples in a particular location. In all cases our method offers improvements over the static display method.

Optimization: We use a very simple gradient descent solver to find the solution. To increase efficiency we subsample all images to 80x40, and in practice this minimization converges fast enough for realtime operation with two viewers. Since this is a research prototype we did not attempt to optimize the code and thus three or more users causes the display to have a slight lag when head positions change.

5. RESULTS

5.1 Crosstalk Reduction

One of the primary issues affecting adoption of autostereoscopic displays is crosstalk. We evaluated our method both visually and quantitatively.

Visual comparisons were made using both imagery intended for multiview devices like ours as well as with test patterns which place a unique numeric image in the corresponding zone. Figure 2 uses photographs to compare the device using static one-to-one mapping with the device using dynamic mapping. Notice that one-to-one mapping produces an image with noticeable crosstalk, while dynamic mapping provides a clean image with little ghosting from other views.

As the number of viewers increases, dynamic mapping will find it harder to place desired images into zones in a way that preserves quality everywhere. Given enough viewers the mapping will revert to one-to-one mapping since this is a good solution for optimizing many eye locations.

5.2 Legacy 2-view content

Most existing content for 3D stereo viewing was captured and stored with only 2-views. Autostereoscopic displays require multiview input. Existing displays often resort to using a fixed pattern for mapping views to zones. This strategy produces some viewing locations which create a correct stereo percept and some locations with stereo inversion.

Another approach is to employ view synthesis methods to produce enough new viewpoints to drive the display.[27,32,33] These methods are compatible with our work since they produce a set of input views. In addition, with dynamic mapping the number of synthesized views need not match the number of device zones, and if eye tracking is available, computation can be reduced by synthesizing only the required viewpoints.

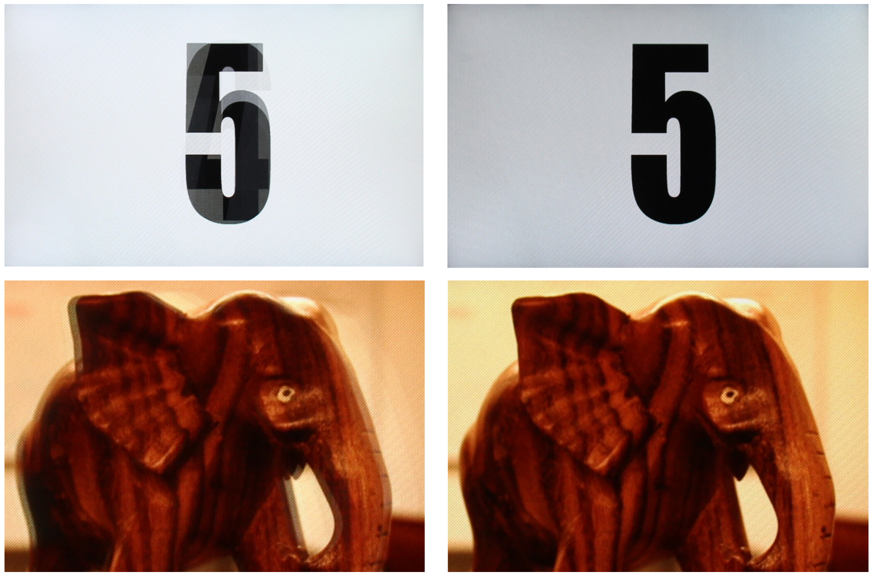

Static Mapping Result   Dynamic Mapping Result

Figure 2. Visual results using static and dynamic mapping are shown. In the numeric test pattern note that static one-to-one mapping has noticeable crosstalk, with multiple numbered views visible. The overall gain of the white background also appears different because of this crosstalk. The elephant example also shows crosstalk in this close up view.

Dynamic mapping also allows 2 channel input to be used directly, without synthesizing additional views, and we concentrate on this case in our discussion below. The desired view at the left eye, wherever it happens to be, is set to the left view provided by the content, $\mathbf{D}(\mathbf{x}_L) = V_L$. This is repeated for the right eye. This allows legacy stereo content to be played on autostereoscopic displays without additional processing steps, which may distort the artistic creator's intention.

The advantage of dynamic mapping in a simple single user case is shown in Figure 3. The best static mapping for 2-view content on our particular display is $\mathbf{Z} = [V_L V_L V_R V_R V_L V_L V_R V_R]$, and thus is used in this example. The observed views for a user in different positions in front of the display are shown allowing the methods to be compared. Notice that with a static mapping the user sometimes sees a correct crosstalk free display, sometimes sees crosstalk, and sometimes sees stereo inversion with the left and right images swapped. This requires users to find the correct location for viewing and then keeping their head still to avoid ruining the effect. In contrast, dynamic mapping allows the user to position themselves in any location and move as needed, since the correct crosstalk free left and right images are always available.

As the number of viewers increases there will eventually be conflicts in which some zones are expected to display both V_L and V_R. These conflicts are unrelated to the displays transfer function, and will not be resolved adequately by the display algorithm's minimization. Consider a perfect multiview display device with no crosstalk. When a single position $\mathbf{x_j}$ is occupied by the left eye of one viewer and the right eye of another we have $\mathbf{D}(\mathbf{x_j}) = V_L$ and $\mathbf{D}(\mathbf{x_j}) = V_R$, thus the least RMSE is achieved by providing a combination of the two images even though this is not visually pleasing. The failure here is in the *view selection policy*, not the display algorithm. A single position should not be expected to display both left and right images.

6. CONCLUSION

In this paper, we propose a simple software-based method to improve the imaging quality of multiview autostereoscopic displays. We have evaluated the method both visually and quantitatively, and found it effective at reducing crosstalk. Since the method dynamically maps views with respect to device specific hardware zones, it can also be used to address other common difficulties such as displaying legacy 2-view content on multiview displays and eliminating stereo inversion.

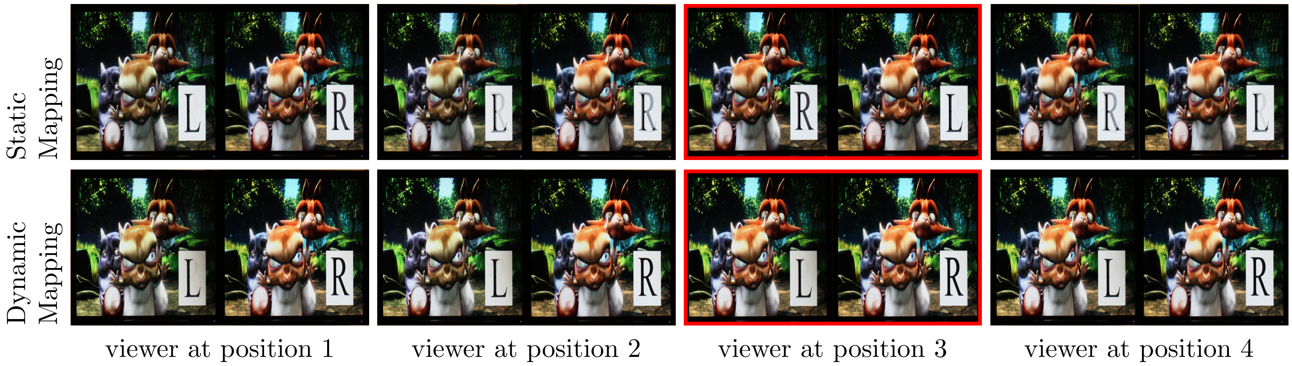

Figure 3. A static LLRRLLRR mapping of legacy 2-view content onto our 8-view display results in stereo inversion for some eye positions, with left and right images swapped. Dynamic mapping provides the correct view to each eye regardless of viewer position.

ACKNOWLEDGMENTS

Partial support from HP, and NSF CCF-0746690.

REFERENCES

[1] Kooi, F. and Toet, A., "Visual comfort of binocular and 3D displays," *Displays* **25**(2-3), 99–108 (2004).

[2] Johnson, R. and Jacobsen, G., "Advances in lenticular lens arrays for visual display," in [*Proceedings of SPIE*], **5874**, 587406 (2005).

[3] Tsai, R., Tsai, C., Lee, K., Wu, C., Lin, L., Huang, K., Hsu, W., Wu, C., Lu, C., Yang, J., et al., "Challenge of 3D LCD displays," in [*Proc. SPIE*], **7329**, 732903–1 (2009).

[4] Matusik, W. and Pfister, H., "3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes," in [*ACM Transactions on Graphics (TOG)*], **23**(3), 814–824, ACM (2004).

[5] Balogh, T., Kovács, P., and Megyesi, Z., "Holovizio 3D display system," in [*Proceedings of the First International Conference on Immersive Telecommunications*], 19, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering) (2007).

[6] Sexton, I. and Surman, P., "Stereoscopic and autostereoscopic display systems," *Signal Processing Magazine, IEEE* **16**(3), 85–99 (1999).

[7] Urey, H., Chellappan, K., Erden, E., and Surman, P., "State of the art in stereoscopic and autostereoscopic displays," *Proceedings of the IEEE* **99**(4), 540–555 (2011).

[8] Holliman, N. S., Dodgson, N. A., Favalora, G. E., and Pockett, L., "Three-dimensional displays: a review and applications analysis," *Broadcasting, IEEE Transactions on* **57**(2), 362–371 (2011).

[9] Wetzstein, G., Lanman, D., Hirsch, M., and Raskar, R., "Tensor Displays: Compressive Light Field Synthesis using Multilayer Displays with Directional Backlighting," *ACM Trans. Graph. (Proc. SIGGRAPH)* **31**(4), 1–11 (2012).

[10] Travis, A., Large, T., Emerton, N., and Bathiche, S., "Wedge optics in flat panel displays," *Proceedings of the IEEE* **101**(1), 45–60 (2013).

[11] Maimone, A., Wetzstein, G., Hirsch, M., Lanman, D., Raskar, R., and Fuchs, H., "Focus 3D: Compressive accommodation display," *ACM Transactions on Graphics (TOG)* **32**(5), 153 (2013).

[12] Redert, A., de Beeck, M., Fehn, C., IJsselsteijn, W., Pollefeys, M., Van Gool, L., Ofek, E., Sexton, I., and Surman, P., "Advanced three-dimensional television system technologies," in [*3D Data Processing Visualization and Transmission, 2002. Proceedings. First International Symposium on*], 313–319, IEEE (2002).

[13] Brar, R., Surman, P., Sexton, I., Bates, R., Lee, W., Hopf, K., Neumann, F., Day, S., and Willman, E., "Laser-based head-tracked 3D display research," *Display Technology, Journal of* **6**(10), 531–543 (2010).

[14] Liou, J., Lee, K., and Huang, J., "Low crosstalk multi-view tracking 3-D display of synchro-signal LED scanning backlight system," *Display Technology, Journal of* **7**(8), 411–419 (2011).

[15] Stolle, H., Olaya, J.-C., Buschbeck, S., Sahm, H., and Schwerdtner, A., "Technical solutions for a full-resolution autostereoscopic 2D/3D display technology," in [*Proceedings of SPIE*], 1–12, International Society for Optics and Photonics (2008).

[16] Kim, S.-K., Yoon, S.-K., and Yoon, K.-H., "Generation of flat viewing zone in dfvz autostereoscopic multi-view 3d display by weighting factor," in [*SPIE Defense, Security, and Sensing*], 87380B–87380B, International Society for Optics and Photonics (2013).

[17] Kim, S.-K., Yoon, S.-K., and Yoon, K.-H., "Crosstalk minimization in autostereoscopic multiveiw 3d display by eye tracking and fusion (overlapping) of viewing zones," in [*SPIE Defense, Security, and Sensing*], 838410–838410, International Society for Optics and Photonics (2012).

[18] Boev, A., Georgiev, M., Gotchev, A., and Egiazarian, K., "Optimized single-viewer mode of multiview autostereoscopic display," in [*Proc. of EUSIPCO*], (2008).

[19] Woodgate, G., Ezra, D., Harrold, J., Holliman, N., Jones, G., and Moseley, R., "Autostereoscopic 3D display systems with observer tracking," *Signal Processing: Image Communication* **14**(1), 131–145 (1998).

[20] Nashel, A. and Fuchs, H., "Random hole display: A non-uniform barrier autostereoscopic display," in [*3DTV Conference: The True Vision-Capture, Transmission and Display of 3D Video, 2009*], 1–4, IEEE (2009).

[21] Ye, G., State, A., and Fuchs, H., "A practical multi-viewer tabletop autostereoscopic display," in [*Mixed and Augmented Reality (ISMAR), 2010 9th IEEE International Symposium on*], 147–156, IEEE (2010).

[22] Heide, F., Wetzstein, G., Raskar, R., and Heidrich, W., "Adaptive image synthesis for compressive displays," *ACM Transactions on Graphics (TOG)* **32**(4), 132 (2013).

[23] Zwicker, M., Matusik, W., Durand, F., and Pfister, H., "Antialiasing for automultiscopic 3D displays," in [*Proceedings of the 17th Eurographics conference on Rendering Techniques*], 73–82, Eurographics Association (2006).

[24] Du, S.-P., Didyk, P., Durand, F., Hu, S.-M., and Matusik, W., "Improving visual quality of view transitions in automultiscopic displays," *ACM Transactions on Graphics (Proceedings SIGGRAPH Asia 2014, Shenzhen, China)* **33**(6) (2014).

[25] Masia, B., Wetzstein, G., Aliaga, C., Raskar, R., and Gutierrez, D., "Display adaptive 3D content remapping," *Computers & Graphics* **37**(8), 983–996 (2013).

[26] Wadhwa, N., Rubinstein, M., Durand, F., and Freeman, W. T., "Phase-based video motion processing," *ACM Transactions on Graphics (TOG)* **32**(4), 80 (2013).

[27] Didyk, P., Sitthi-Amorn, P., Freeman, W., Durand, F., and Matusik, W., "Joint view expansion and filtering for automultiscopic 3D displays," *ACM Transactions on Graphics (TOG)* **32**(6), 221 (2013).

[28] Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P., "Image quality assessment: From error visibility to structural similarity," *Image Processing, IEEE Transactions on* **13**(4), 600–612 (2004).

[29] Dodgson, N. A., "Variation and extrema of human interpupillary distance," in [*Proceedings of SPIE*], **5291**, 36–46 (2004).

[30] Hansen, D. W. and Ji, Q., "In the eye of the beholder: A survey of models for eyes and gaze," *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **32**(3), 478–500 (2010).

[31] Viola, P. and Jones, M., "Rapid object detection using a boosted cascade of simple features," in [*Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*], **1**, I–511, IEEE (2001).

[32] Smolic, A., Muller, K., Dix, K., Merkle, P., Kauff, P., and Wiegand, T., "Intermediate view interpolation based on multiview video plus depth for advanced 3D video systems," in [*Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*], 2448–2451, IEEE (2008).

[33] Farre, M., Wang, O., Lang, M., Stefanoski, N., Hornung, A., and Smolic, A., "Automatic content creation for multiview autostereoscopic displays using image domain warping," in [*Multimedia and Expo (ICME), 2011 IEEE International Conference on*], 1–6, IEEE (2011).